

CLAIMS

1. A method of optimizing traffic in an internetwork, the method comprising:

5 selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

10 measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics including one or more of:

launching a first plurality of one or more packets, and measuring the first plurality of one or more packets;

15 measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched;

monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets;

20 serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows; and

25 encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

30 calculating at least one performance metric for the subset of flows in the at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and

in response to calculating the at least one performance metric,
rearchitecting the internetwork to optimize one or more of the at least one
performance metric, rearchitecting the internetwork including at least one of:

altering a plurality of one or more routing tables in the internetwork,
5 wherein the plurality of one or more routing tables include at least one of:
network-layer routing tables, layer 3 routing tables, IP routing tables, layer 2
forwarding tables, and MPLS forwarding tables;

redirecting the subset of flows to a second internetwork coupled
to the internetwork at one or more Points of Presence; and

10 affecting forwarding decisions of the subset of flows, by
imposing one or more of: NAT, GRE, and tunneling techniques other
than GRE.

2. The method of claim 1, wherein the plurality of one or more routing
15 tables includes network-layer routing tables.

3. The method of claim 1, wherein the plurality of one or more routing
tables includes layer 3 routing tables.

20 4. The method of claim 1, wherein the plurality of one or more routing
tables includes IP routing tables.

5. The method of claim 1, wherein the plurality of one or more routing
tables includes layer 2 routing tables.

25 6. The method of claim 1, wherein the plurality of one or more routing
tables includes MPLS forwarding tables.

7. The method of claim 1, wherein the performance characteristics includes
30 the plurality of one or more round trip measurements.

8. The method of claim 1, wherein the performance characteristics includes
the plurality of one or more one-way measurements

9. The method of claim 1, wherein the one or more of the at least one performance metric to be optimized quantifies path performance of the subset of flows.

5

10. The method of claim 9, wherein the one or more of the at least one performance metric characterizes a quality of a network application.

11. The method of claim 10, wherein the application includes http.

10

12. The method of claim 10, wherein the application includes ftp.

13. The method of claim 10, wherein the application includes telnet.

15

14. The method of claim 10, wherein the application includes voice.

15. The method of claim 10, wherein the application include video.

20

16. The method of claim 1, wherein one or more of the at least one performance metric quantifies one or more of: overall usage of the at least the portion of the internetwork, absolute individual network link usage in the at least the portion of the internetwork, relative individual network link usage in the at least the portion of the internetwork, and link usage cost in the at least the portion of the internetwork.

25

17. The method of claim 1, wherein the subset of flows include at least a first sub-plurality of one or more flows, and a second sub-plurality of one or more flows, and the first sub-plurality of one or more flows uses at least a first measured performance characteristic, and the second sub-plurality of one or more flows uses at least a second measured performance characteristic.

30

18. The method of claim 1, wherein the one or more of the at least one performance metric quantifies geographic distance covered by the subset of flows in the internetwork.
19. The method of claim 1, wherein the selection of the subset of flows is user-based in the at least the portion of the internetwork.
20. The method of claim 1, wherein the internetwork comprises an autonomous sub-system of a larger network.
21. The method of claim 20, wherein the larger network is the Internet.
22. The method of claim 20, wherein the internetwork is a BGP autonomous system.
23. The method of claim 1, wherein the internetwork comprises an autonomous system of a larger network.
24. The method of claim 23, wherein the larger network is the Internet.
25. The method of claim 23, wherein the internetwork is a BGP autonomous system.
26. The method of claim 1, wherein the internetwork comprises an overlay network.
27. The method of claim 1, wherein the internetwork comprises a plurality of one or more networks, the plurality of one or more networks coupled together, wherein the plurality of one or more networks include one or more of local-area networks, metropolitan-area networks, and wide-area networks.
28. The method of claim 1, wherein the measuring performance characteristics includes launching a first plurality of one or more packets, and measuring the first plurality of one or more packets.

29. The method of claim 28, wherein the measuring performance characteristics includes launching a first plurality of one or more round-trip packets, and measuring the first plurality of one or more round-trip packets, such that the plurality of one or more round trip measurements include round trip measurements for the launched packets.
30. The method of claim 29, wherein the first plurality of one or more round-trip packets include traceroute ICMP packets.
31. The method of claim 29, wherein the first plurality of one or more round-trip packets include ping ICMP packets.
32. The method of claim 29, wherein the first plurality of one or more round-trip packets include telnet packets.
33. The method of claim 29, wherein the first plurality of one or more round-trip packets include TCP packets from an empty TCP transaction.
34. The method of claim 29, wherein the first plurality of one or more round-trip packets include http packets.
35. The method of claim 28, wherein the measuring performance characteristics includes launching a first plurality of one or more one-way packets, and measuring the first plurality of one or more one-way packets, such that the plurality of one or more one-way measurements include one-way measurements for the launched packets.
36. The method of 1, wherein the measuring performance characteristics includes measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched

37. The method of claim 1, wherein round-trip measurements include one or more of round-trip delay, round-trip jitter, round-trip loss, round-trip available bandwidth, and round-trip total bandwidth.

5

38. The method of claim 37, wherein round-trip measurements include round-trip delay.

39. The method of claim 37, wherein round-trip measurements include round-trip jitter.

10

40. The method of claim 37, wherein round-trip measurements include round-trip loss.

41. The method of claim 37, wherein round-trip measurements include round-trip available bandwidth.

15

42. The method of claim 37, wherein round-trip measurements include round-trip total bandwidth.

20

43. The method of claim 1, wherein one-way measurements include one or more of one-way delay, one-way jitter, one-way loss, one-way available bandwidth, and one-way total bandwidth

44. The method of claim 43, wherein one-way measurements include one-way delay.

25

45. The method of claim 43, wherein one-way measurements include one-way jitter.

30

46. The method of claim 43, wherein one-way measurements include one-way loss.

47. The method of claim 43, wherein one-way measurements include one-way available bandwidth.

48. The method of claim 43, wherein one-way measurements include one-way total bandwidth.

49. The method of claim 1, wherein altering a plurality of one or more routing tables in the internetwork, wherein the plurality of one or more routing tables include at least one of: network-layer routing tables, layer 3 routing tables, IP routing tables, layer 2 forwarding tables, and MPLS forwarding tables.

50. The method of claim 49, wherein the altering of the plurality of one or more routing tables is applied automatically.

51. The method of claim 49, wherein the altering of the plurality of one or more routing tables includes configuring a plurality of one or more routers , wherein the configuring the plurality of one or more routers statically alter the routing of flows .

52. The method of claim 51, wherein the configuring a plurality of one or more routers includes route maps.

53. The method of claim 51, wherein the configuring a plurality of one or more routers includes static route statements.

54. The method of claim 49, wherein the altering of the plurality of one or more routing tables includes configuring a plurality of one or more routers , wherein the configuring the plurality of one or more routers adjust the processing of dynamic routing updates.

55. The method of claim 49, wherein the altering of the plurality of one or more routing tables includes configuring a plurality of one or more routers ,

wherein the configuring the plurality of one or more routers adjust the processing of dynamic routing updates.

56. The method of claim 49, wherein the altering of the plurality of one or more routing tables is performed at least partly by a plurality of one or more dynamic routing protocols.

57. The method of claim 56, wherein the plurality of one or more dynamic routing protocols includes BGP.

58. The method of claim 56, wherein the plurality of one or more dynamic routing protocols includes a plurality of one or more IGP routing protocols.

59. The method of claim 58, wherein the plurality of one or more IGP routing protocols includes OSPF.

60. The method of claim 49, wherein the altering of the plurality of one or more routing tables are applied manually by a user.

61. The method of claim 1, wherein the rearchitecting of the internetwork includes redirecting the subset of flows to a second internetwork coupled to the internetwork at one or more Points of Presence.

62. The method of claim 61, wherein the redirecting of the subset of flows is at least partly across an exit point traversed by at least a portion of the subset of flows from the internetwork, wherein the at least one performance metric is optimized.

63. The method of claim 1, wherein the optimizing of the performance metric includes:
calculating at least one performance metric for a plurality of one of more paths in the at least the portion of the internetwork; and

at least partly responsive to the calculating of the at least one performance metric, selecting a path in the at least the portion of the internetwork from the plurality of one or more paths.

5 64. The method of claim 63, wherein the selected path is a direct path.

65. The method of claim 63, wherein the selected path is an indirect path.

10 66. The method of claim 1, wherein the optimizing of the performance metric includes:
selecting a plurality of one or more exit points; and
selecting a plurality of one or more paths to reach the plurality of one or more exit points in the at least the portion of the internetwork.

15 67. The method of claim 66, wherein the steps of selecting the plurality of one or more exit points and selecting the plurality of one or more paths are performed separately

20 68. The method of claim 66, wherein the steps of selecting the plurality of one or more exit points and selecting the plurality of one or more paths are performed together

25 69. The method of claim 1, wherein the measuring performance characteristics includes monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets.

30 70. The method of claim 1, wherein the measuring performance characteristics includes serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows.

71. The method of claim 1, wherein the measuring performance characteristics includes encoding material within a third plurality of one or more

flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows.

72. The method of claim 1, wherein the measuring performance characteristics is done at least partly using flow information export.

73. The method of claim 1, wherein the measuring performance characteristics is done using RMON II.

74. The method of claim 1, wherein the measuring performance characteristics is provided by a source external to the subset of flows.

75. A method of optimizing traffic in an internetwork, the method comprising:
selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;
measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics includes one or more of:
measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched;
monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets;

5 serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows; and encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

10 calculating at least one performance metric for the subset of flows in the at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and in response to calculating the at least one performance metric, affecting the routing of the subset of flows by altering a plurality of one or more DNS entries in the internetwork.

76. A network systems, comprising:

20 a plurality of one or more network devices configured such that, when the plurality of one more network devices is deployed in an internetwork, the plurality of one or more network devices performs:

selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

25 measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics including one or more of:

30 launching a first plurality of one or more packets, and measuring the first plurality of one or more packets;

measuring a second plurality of one or more packets in the
internetwork, wherein the second plurality of one or more packets were
already launched;

5 monitoring a first plurality of one or more flows in the
internetwork, duplicating one or more packets from the plurality of one
or more flows, and measuring the one or more duplicated packets;

serving as a proxy hop for a second plurality of one or more
flows, and measuring the second plurality of one or more flows; and

10 encoding material within a third plurality of one or more flows, causing
a fourth plurality of one or more flows to be generated, wherein the third
plurality of one or more flows traverses a first path of the internetwork, and the
fourth plurality of one or more flows traverses a second path of the
internetwork, wherein at least a portion of the first path of the internetwork and
at least a portion of the second path of the internetwork are equal, and
15 measuring a subset of the fourth plurality of one or more flows;

calculating at least one performance metric for the subset of flows in the
at least a portion of the internetwork, the at least one performance metric at least
partly determined from the measured performance characteristics; and

20 in response to calculating the at least one performance metric,
rearchitecting the internetwork to optimize one or more of the at least one
performance metric, rearchitecting the internetwork including at least one of:

25 altering a plurality of one or more routing tables in the internetwork,
wherein the plurality of one or more routing tables include at least one of:
network-layer routing tables, layer 3 routing tables, IP routing tables, layer 2
forwarding tables, and MPLS forwarding tables;

redirecting the subset of flows to a second internetwork coupled
to the internetwork at one or more Points of Presence; and

30 affecting forwarding decisions of the subset of flows, by
imposing one or more of: NAT, GRE, and tunneling techniques other
than GRE.

77. The network system of claim 76, where the rearchitecting of the internetwork is performed within the device.

78. A network systems, comprising:

5 a plurality of one or more network devices configured such that, when the plurality of one more network devices is deployed in an internetwork, the plurality of one or more network devices performs:

selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

10 measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics including one
15 or more of:

measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched;

20 monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets;

serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows; and

25 encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the
30 internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

calculating at least one performance metric for the subset of flows in the at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and

5 in response to calculating the at least one performance metric, affecting the routing of the subset of flows by altering a plurality of one or more DNS entries in the internetwork

79. The network system of claim 78, where the affecting the routing of the subset of flows is performed within the device.

10

80. A method of populating a plurality of one or more groups with a plurality of one or more network addresses, the method comprising:

15 selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity;

creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices;

20 populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

25 including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

81. The method of claim 80, wherein at least path performance is selected.

82. The method of claim 80, wherein at least network cost is selected.

30

83. The method of claim 80, wherein at least network usage is selected.

84. The method of claim 80, wherein at least geographical proximity is selected.
85. The method of claim 80, wherein at least topological proximity is selected.
86. The method of claim 80, wherein at least statistical similarity is selected.
87. The method of claim 80, wherein the at least the classification function includes a plurality of one or more criteria, the plurality of one or more criteria including at least one of topological proximity and administrative policy, wherein the topological proximity includes one or more of adjacency and dependency, and the administrative policy includes one or more of address aggregation and user-defined policy.
88. The method of claim 87, wherein the plurality of one or more criteria includes at least topological proximity, wherein the topological proximity includes one or more of adjacency and dependency.
89. The method of claim 87, wherein the plurality of one or more criteria includes at least administrative policy, wherein the administrative policy includes one or more of address aggregation and user-defined policy.
90. The method of claim 80, wherein, after the populating, the plurality of one or more groups includes a group, the group including network addresses corresponding to the plurality of one or more metrics having a value in a pre-defined range.
91. The method of claim 80, wherein, after the populating, the plurality of one or more groups includes clusters created by a clustering algorithm applied to at least one of the plurality of one or more metrics.

92. The method of claim 80, wherein the selecting, the creating, and the classifying are performed at a central server in an autonomous system.

5 93. The method of claim 80, wherein the selecting, the creating, and the classifying are performed at a plurality of one or more distributed servers distributed through a plurality of one or more autonomous systems.

10 94. The method of claim 80, wherein the selecting, creating, and populating are repeated, at least partly responsive to one or more of: a change in the at least one of the plurality of one or more metrics, an expiration of a time period, and an external input.

15 95. The method of 94, wherein the selecting, creating, and populating are repeated, at least partly responsive to the change in the at least one of the one or more metrics.

96. The method of 94, wherein the selecting, creating, and populating are repeated, at least partly responsive to the expiration of a time period.

20 97. The method of 94, wherein the selecting, creating, and populating are repeated, at least partly responsive to the external input.

25 98. The method of 94, wherein, prior to the change, a first plurality of one or more groups is populated, and after the change, a second plurality of one or more groups is populated, and the first plurality of one or more groups is at least partly different from the second plurality of one or more groups.

30 99. The method of claim 80, further comprising:
prior to populating the plurality of groups, creating a default population for at least one of the plurality of one or more groups, wherein the default population is derived at least partly from an observed routing table.

100. The method of claim 99, wherein the observed routing table is from an autonomous system, such that each of the plurality of network addresses is reachable by the autonomous system.

5 101. The method of claim 99, wherein the default population is a subset of the observed routing are performed on an existing routing table.

10 102. The method of 101, wherein after selecting, creating, and populating are performed on the existing routing table, any changes to the existing routing table are limited to combining two or more entries into one group.

15 103. The method of 101, wherein after selecting, creating, and populating are performed on the existing routing table, changes to the existing routing table include dividing at least one entry into two or more groups.

20 104. The method of claim 80, further comprising:
prior to populating the plurality of groups, creating a default population for at least one of the plurality of one or more groups, wherein the default population is derived at least partly from observed network traffic.

25 105. The method of claim 80, wherein the selecting, creating, and populating are performed in a first autonomous system.

30 106. The method of claim 105, further comprising:
prior to populating the plurality of groups, creating a default population for at least one of the plurality of one or more groups, wherein the default population is derived at least partly from an observed routing table.

35 107. The method of 105, wherein the observed routing table is from the first autonomous system.

108. The method of 105, wherein the observed routing table is from a second autonomous system.

109. The method of claim 80, wherein at least one subset of the plurality of one or more groups includes contiguous addresses.

5 110. The method of claim 80, wherein at least one subset of the plurality of one or more groups includes at least two discontinuous addresses.

111. The method of claim 80, where the plurality of one or more network addresses includes one or more IP addresses.

10

112. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

15

responsive to 1) selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity, and 2) creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices,

20

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

25

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

30 113. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing

devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity;

creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices,

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

114. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

responsive to selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity,

creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices;

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

5 115. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

10 selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity;

responsive to creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices,

15 populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

20 including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.